

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant: Athena CHRISTODOULOU et al.

Title: ESTABLISHMENT OF NETWORK
CONNECTIONS

Appl. No.: Unassigned

Filing Date: July 30, 2003

Examiner: Unassigned

Art Unit: Unassigned

CLAIM FOR CONVENTION PRIORITY

Commissioner for Patents
PO Box 1450
Alexandria, Virginia 22313-1450

Sir:

The benefit of the filing dates of the following prior foreign applications filed in the following foreign country is hereby requested, and the right of priority provided in 35 U.S.C. § 119 is hereby claimed.

In support of this claim, filed herewith are certified copies of said original foreign applications:


- GREAT BRITAIN Patent Application No. 0217795.4 filed 07/31/2002.
- GREAT BRITAIN Patent Application No. 0222696.7 filed 10/01/2002.

Respectfully submitted,

Date July 30, 2003

HEWLETT-PACKARD COMPANY
Intellectual Property Administration
P.O. Box 272400
Fort Collins, Colorado 80527-2400

By



William T. Ellis
Attorney for Applicant
Registration No. 26,874

100



INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed


Dated 3 July 2003



01AUG02 E737683-1 001463
P01/7700 0.00-0217795.4

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)



The Patent Office

Cardiff Road
Newport
South Wales
NP10 8QQ

1. Your reference

2. 0217795.4

31 JUL 2002

3. Full name, address and postcode of the or of each applicant (*underline all surnames*)

Hewlett-Packard Company
3000 Hanover Street
Palo Alto
CA 94304, USA

Patents ADP number (*if you know it*)

Delaware, USA

If the applicant is a corporate body, give the country/state of its incorporation

496588001

4. Title of the invention Establishment of Network Connections

5. Name of your agent (*if you have one*)

Bruce G R Jones
Hewlett-Packard Ltd, IP Section
Filton Road, Stoke Gifford
Bristol BS34 8QZ

"Address for service" in the United Kingdom to which all correspondence should be sent (*including the postcode*)

8072258001

Patents ADP number (*if you know it*)

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (*if you know it*) the or each application number

Country

Priority application number
(*if you know it*)

Date of filing
(*day / month / year*)

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application

Date of filing
(*day / month / year*)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (*Answer 'Yes' if:*

Yes

- a) any applicant named in part 3 is not an inventor, or
 - b) there is an inventor who is not named as an applicant, or
 - c) any named applicant is a corporate body.
- See note (d))

Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form

Description	10
Claim(s)	2
Abstract	1
Drawing(s)	9 + 9

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (*Patents Form 7/77*)

Request for preliminary examination and search (*Patents Form 9/77*)

1

Request for substantive examination (*Patents Form 10/77*)

Any other documents
(please specify)

Fee Sheet

11. I/We request the grant of a patent on the basis of this application.

Signature

Date

Bruce Graeme Roland Jones 31 July 2002

12. Name and daytime telephone number of person to contact in the United Kingdom

K Nommeots-Nomm Tel: 0117-312-9947

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- a) If you need help to fill in this form or you have any questions, please contact the Patent Office on 08459 500505.
- b) Write your answers in capital letters using black ink or you may type them.
- c) If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- d) If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- e) Once you have filled in the form you must remember to sign and date it.
- f) For details of the fee and ways to pay please contact the Patent Office.

ESTABLISHMENT OF NETWORK CONNECTIONS

The present invention relates to the establishment of network connections, such as for example, the establishment of a connection from a client computing entity ("client")
5 to a server computing entity ("server") which hosts media content located at a level in the hierarchical architecture of a website below the primary or "home" page for that site.

The home page of a website is often simply a series of pointers to other parts of the
10 site (or indeed to other distinct sites, which in the context of the present application are nonetheless regarded as being "below" the home page in a hierarchy because they are reachable via a link on the home page). These pointers are usually implemented by one or more hyperlinks, and so for new visitors to the site availability of the home page is important if they are to be able easily to navigate the site (and where
15 appropriate, e.g. in circumstances outlined above, any associated site) to the fullest extent possible. When a connection is made to a particular web site, the initial connection is therefore usually established with a primary or main server because it hosts the home page, and results in the main server returning a copy of the home page to the browser programme within the requesting client. In practice several main
20 servers are likely to be employed, with one main server being the master and the other main servers being slaves to that master. This arrangement enables content changes which are implemented on the master main server to be automatically replicated on each of the slaves. Increasingly, the provision of constant availability and consistently high performance of a website are seen as important. Therefore, because
25 the majority of all new network traffic to the site will initially be requesting the home page, and will therefore be directed to one of the main servers hosting the home page, maintenance of the main servers' ability to provide such availability and performance is regarded as critical, which in turn means that any measures which can be taken to reduce load on the main servers are potentially valuable.

30

One commonly employed such measure is to host pages which are accessible from the home page via the actuation of the links thereon on one or more further servers separate to the main server or servers, known in this application as secondary servers. This has several advantages: firstly the load on the main server(s) is, comparatively

speaking, reduced, since when a link on the home page is actuated, the sub-page (so called because it is located at a lower level in the hierarchy of the website architecture than the home page) to which that link points will be located on one of the secondary servers. Thus upon actuation of a link to a sub-page the browser within the requesting client will be connected to the relevant secondary server by the main server, so that the secondary server and not the main server will be performing all of the relevant tasks in accordance with hypertext transfer protocol (http), and will return a copy of the sub-page to the browser within the requesting client. This means that, even at times of extremely heavy traffic, the main server is simply dealing with requests for the home page, and possibly also passing requests for sub-pages to the secondary server or servers, rather than actually processing requests for sub-pages, which are frequently richer in content than the home page and therefore more apt to require greater time to transfer from a server to a client.

A second advantage of this approach is that it enables the provision of a degree of fault tolerance to failure or overloading of a secondary server. If several secondary servers are in service, and there is at least some degree of duplication in the various sub-pages that they host, a request for a particular sub-page may be directed to any secondary server on which that page is hosted, thus reducing the possibility that any sub-page is unavailable. This process may be performed simply within the main server on the basis of contemporaneous circumstances, such as load on the various secondary servers or their status (i.e. whether they are fully operational or not), since it is a relatively rapid process to redirect a request for a sub-page to a given secondary server. Alternatively, this advantage may be realized during the process of converting or "resolving" the alpha-numeric characters used by people to identify a website (e.g. www.bbc.co.uk), usually known as a Website name or a Uniform Resource Locator ("URL"), into an Internet Protocol address ("IP address") identifying the address of a particular server within the Internet (an example of which is 192.168.45.4). This process is performed by a Domain Name Service ("DNS") server, and in accordance with this process, a given sub-page may simultaneously be hosted on several secondary servers, all of which have different IP addresses. Actuation of a link to request such a sub-page may be resolved to the IP addresses of any one of these secondary servers at the DNS server on the basis of contemporaneous circumstances, and may be simply on the basis of a count maintained within the DNS server

regarding the number of connection requests each secondary server has received, with the aim of maintaining at least a degree of equality in the numbers of connections to each secondary server and thereby providing balancing of load on the various secondary servers. Alternatively, in the event that one server is faulty,
5 resolution to an IP address of a particular secondary server may be performed to direct all requests for the relevant links to one or more other secondary servers.

A first aspect of the present invention provides an alternative manner of realising advantages of using one or more secondary servers, in which a web page (typically,
10 but not necessarily a home page) sent to a client has embedded within it one or more links to a sub-page, and one or more of these links contain predetermined signifiers indicating a particular secondary server from which the sub-page is to be retrieved. These predetermined signifiers could for example be specific characters within the URL or the whole of a URL for a link to a particular sub-page which are then
15 resolved, either at a DNS server or at the main server, to an IP address of a given secondary server. Alternatively they may simply be the IP address of a given secondary server thus avoiding the need for subsequent resolution.

By providing predetermined signifiers within a link for the sub-page, the task of
20 deciding which secondary server a request for a sub-page should be directed to upon actuation of the link is devolved to the client, thereby further reducing the load on the main server or servers, or indeed a potential bottleneck at the DNS server while this decision is made.

In a preferred embodiment a link to a sub-page is actually an alias for a plurality of alternate links each of which has a signifier corresponding to the network addresses of one of a plurality of secondary servers. In this manner a degree of tolerance to faults within a particular secondary server is provided; if one secondary server is not operational or excessively slow, then actuation of an alternate link will attempt to
25 achieve connection to another secondary server. The plurality of alternate links may be configured to actuate sequentially, i.e. upon failure of a first alternate link (e.g. as defined by timeout) a second alternate link is automatically actuated, and so on.
30 Alternatively, plural alternate links may be configured to actuate simultaneously, with the first successful establishment of a connection to a secondary server causing all

other alternate attempts to connect to other secondary servers to abort. In such an instance it is preferable for the secondary servers to be provided with a mechanism for denying what may be defined as abusive use of simultaneous connection attempts, which may be done by identifying the browser programme (from data contained
 5 within the http request) and in the event of an abuse being identified, denying one or more of the simultaneous connection.

Embodiments of the invention will now be described, by way of example, and with reference to the accompanying drawings in which:

10

Figs. 1 to 8 are schematic illustrations of the process of establishing connections between one or more clients and servers for the purpose of retrieving web pages; and

Fig. 9 is a listing of javascript associated with multiple alternate links.

15

Referring now to Fig. 1, first and second client computing entities 10, 12 are connected to the Internet. The web browsing programme (not illustrated specifically) of the first client 10 is seeking to connect to a website (fictitious - at the time of writing) providing safari information, whose URL is <http://www.safari.fun.co.uk>, and
 20 which is hosted on a primary or main server 20. In practice connecting to this website actually means downloading a copy of its home page from the main server 20, which for websites supporting heavy traffic, is likely to be one of a plurality of main servers, one of which is a master to which the others are slaved. Such an arrangement provides the power of several servers to support heavy traffic for the home page, but
 25 at the same time means that in the event that the content on the home page is to be changed, it only need be changed on the master main server, whereupon the content on each slave main server has its content reconciled with that on the master in a manner known *per se*, and which will therefore not be discussed further.

30 In order to connect to a main server (only one of which is illustrated herein) 20, the URL of the website first needs to be converted, or "resolved" into an Internet Protocol address ("IP address"), which is a series of numbers signifying the location within the internet of the server to which it is desired to establish connection. This process takes place at what is known as a Domain Name Service ("DNS") server 30, whose own IP

address will typically be stored in the client 10 as part of the process of connecting the client 10 to the Internet. In Fig. 1 the process of the browser connecting to the DNS server 30 and the resolution of the URL for the requested website into the IP address of the main server 20 by the DNS server 30 are illustrated schematically. Referring now to Fig. 2, once connection of the client 10 and the main server 20 has been established, the main server returns a copy of the home page 100 to the client 10. It can be seen from Fig. 2 that the home page contains two links: one for "hunting" 102, and the other for "wildlife" 104. In the present application the term link is intended to include within its scope a pointer from one location to another, which is actuatable to cause connection from the location of the link to the location to which the link points. In the present example these links are hyperlinks, that is to say either of these may be actuated by clicking upon the relevant icon to navigate to the particular page at which their subject matter is located. Such a page is known hereinafter as a sub-page, because in each case it lies below the home page in the architectural hierarchy of the website. (NB The use of the term hierarchy in the present application is intended to apply broadly, so that for example a first web page is denoted herein as lying "below" a second web page in a hierarchy if it is accessible via a link on the second page, even though it may for example be a page on an entirely independent site.) Actuation of the link 104 causes the browser programme within the client 10 to seek connection to a predetermined IP address at which the subject matter of that link is located. This IP address is usually coded in terms of a URL, such as in the present case:

<http://www.safarifun/wildlife.html>

although in the present example this is really an alias for another URL which is the URL actually incorporated within the link:

<http://www.safarifun/wildlifeone.html>

and which is resolvable to the IP address of a particular secondary server on which the sub-page "wildlife" is hosted. An alias is typically used to prevent a user becoming aware of the existence of the actual URL for the secondary server, and one reason for this is that different URLs for different secondary servers hosting the same wildlife sub-page may be provided to different clients under the same alias. An example of

this is illustrated in Fig. 3, where the client 12 has a different copy of a home page which has the same alias URL for the wildlife link, but a different actual URL, ending in ".../wildlifetwo.html", and which therefore identifies a different secondary server to the link on the page provided to the first client 10.

5

Referring now to Fig. 4, when the wildlife link 104 is actuated (this action being signified by the "action" graphic around the link 104) in the browser of the first client 10, the browser connects to the DNS server to obtain resolution of the URL into an IP address. In this first example, the DNS server is able only to resolve the primary URL, that is to say: http://www.safarifun.co.uk, because it does not "recognize" the subsequent character string (i.e. the website administrator has not registered an IP address for the URL as a whole with the DNS). Therefore, in the first instance actuation of the link 104 causes connection to the DNS and resolution of the primary URL to the IP address of the main server 20. Once connection to the main server 20 has been established, the main server 20 resolves the full URL to an IP address of a secondary server 201 at which the sub-page "wildlife" is hosted. This resolution is performed in accordance with information within a look-up table typically stored in the memory of the main server 20, and which was established at the time the home page containing the link 104 was sent out to the first client 10. Following resolution of the full URL to the IP address of the secondary server 201, the main server redirects connection of the client to the secondary server 201.

A significant distinction between the sequence of events as set out above and the prior art is that whereas in the present invention, the link 104 sent with the home page to client 10 contains a signifier (which in this instance is the character string: "/wildlifeone.html") within the URL which is resolved by the main server to the IP address of secondary server 201 in accordance with a lookup table, in the prior art, where such a resolution is performed by the main server to divert a request for a sub-page to a secondary server, this is done dynamically, i.e. on the basis of a decision made contemporaneously. Thus in the above-described embodiment, under normal operation, no decision-making process takes place at the main server with regard to the destination secondary server.

Referring now to Fig. 5, the link 104 which is sent with a copy of the home page to the second client 12 is actually different to that sent to the first client. When the link 104 is actuated within the browser of the second client 12, the URL for which resolution is requested at the DNS server 30 is therefore correspondingly different to that requested in connection with actuation of the link 104 in the browser of the first client, with the characters subsequent to the primary URL (i.e. those characters which are indicative of the sub-page being requested) in the instance of the second client being “../wildlifetwo.html”, as opposed to “../wildlifeone.html” in the instance of the first client 10. As in the case of the scenario of Fig. 4, the primary part of the URL is resolved to the IP address of the main server 20, which then resolves the character string “wildlifetwo.html” using the lookup table to the IP address of the secondary server 202, where another copy of the sub-page is located, and the main server 20 then passes connection of the second client to the secondary server 202. By sending home pages having different URLs for what is ostensibly the same link to a sub-page, it is possible to balance the load on the differing secondary servers, and this is one reason why these links are aliased, i.e. when actuated, the URL shown in the address bar of the browser is not the URL for which resolution is obtained at the main server 20.

Referring now to Fig. 6, in accordance with a second embodiment, the full URL of the link provided to each of the clients 10, 12 is resolvable at the DNS server 30 to an IP address. Thus, in the scenario of Fig. 6, the URL of the link 102 actuated by the browser of the first client 10 is resolved to the IP address of the secondary server 201, and connection is then established directly with the secondary server 201 without first passing to the main server 20. In accordance with this modified embodiment a similar scenario occurs when the link 102 is actuated by the browser of the second client, with the URL of that link resolving at the DNS to the IP address of the secondary server 202, but this has not been illustrated in a separate figure for brevity's sake.

Fig. 7 illustrates yet a further alternative, in which the link 104 is provided with the home page to the browser in the form of an IP address of the secondary server 201 at which the sub-page “wildlife” is hosted. This has the advantage that it obviates the need for resolution of a URL to an IP address, but the corresponding disadvantage that if it is desired to alter the IP address of the secondary server 201, this link will

fail, whereas if a URL is used, any change in the IP address of the secondary server (provided that this change is recorded with the DNS server 30) will have no effect upon the connection to the secondary server. Again, a corresponding scenario involving the actuation of the link 104 in the browser of the second client 12 has not
 5 been illustrated since it adds nothing to the understanding of the scenario.

The various embodiments of the invention thus far described illustrate the principle of providing, in the link sent with a home page, a signifier (whether this is in the form of a particular URL, or an IP address) indicative of the destination secondary server, so
 10 that the issue of which secondary server to which a client is to be connected upon actuation of the link does not need to be dealt with by a main server. There are however further advantages of providing such links. Referring now to Fig. 8, in a modification of the scenarios previously described, when the main server sends a home page to a client, the home page includes plural links to the wildlife sub-page.
 15 As previously these links all have an alias, for the reasons previously described, and in this case, all of the plural links have the same alias. When for example the wildlife link 104 is actuated, should connection to the secondary server identified by the first listed URL fail to be established within a given time period, machine-executable code associated with the link, which in the present example has the form of javascript (this
 20 being illustrated in Fig. 9), abort this connection attempt, and then attempt connection to the secondary server identified by the second URL in the list. The machine-executable code associated with the link thus provides for sequential attempts to connect to the different secondary servers identified by the different URLs in the event of a timeout failure to connect to any one of the listed URLs.

25

A further aspect of this modification is that, in the event that the same set of alternate links are provided to each requesting client, the links are preferably provided in a variety of orders so that the load on the various secondary servers corresponding to the different URLs is at least approximately balanced. Thus in the illustrated example of
 30 Fig. 8, the first client 201 and the second client 202 have received the same four alternate links but in a different order. One manner in which this may be achieved is simply to provide the links in a random order on each occasion, which will therefore, for large numbers, ensure an approximately equal distribution of loading liability for each secondary server. Alternatively, the order in which the links are provided could

be monitored continually with a log being kept of the various occasions a given alternate link has been provided at a given place in the order, and the log used to ensure equal distribution of loading liability.

In a further modification, an even greater number of alternate links may be provided, and by analogy with the example of Figs. 3 to 7, different sets of alternate links are sent to different clients, further balancing the load on the corresponding secondary servers.

The sequential use of alternate links in the manners described above provides tolerance at the client side of faults at the server side. However, sequential actuation can be time consuming, and if speed of connection is an important parameter then in a further modification, it is possible to configure the code associated with the links to cause two or more, and preferably all of the alternate links to activate simultaneously. This provides the advantage to the client that the fastest performing link on any given occasion will always establish a connection within the shortest possible time, which is not necessarily the case with sequential actuation if the fastest link is not the first-actuated link. Preferably, in order to avoid excessive duplication, the connections sought by the slower links will be aborted at some predetermined milestone in the course of the establishment of a full connection by the fastest link, for example upon having found the sub-page, for example.

The possibility for simultaneous actuation of one or more alternate links however potentially creates a problem for the secondary servers, since such a mode of operation is open to abuse, with the result that it generates a substantial amount of redundant load upon the secondary servers as a result of all of the connections which are sought and then aborted, and therefore potentially damages attempts to balance load upon the secondary servers. Indeed these potentially damaging consequences are a potential outcome quite generally whenever a browser in a client does not follow the actions set out in the code associated with whatever links have been provided to it. In order to provide the possibility of reducing damage caused by deviant client browser behaviour, data relating to interaction behaviour of a given browser (which identifies itself to a server each time an http request is made) in a server log can be used to match the actual behaviour of a given client browser to the anticipated behaviour of that client browser based on the nature of the alternate links and associated code

issued to that browser. If significant deviation is found between the two then an assumption can be made to the effect that the client browser has been hacked to adopt behaviour deviant to that intended, and the client browser's access to one or more of the secondary servers can be reduced or removed as appropriate.

CLAIMS

1. A method of providing a sub-page of a website to a requesting client comprising the steps of:
 - 5 sending to the client a copy of a first web page which includes at least one link which is actuatable to connect the client to a sub-page; wherein:
 - the at least one link comprises at least one signifier identifying at least one predetermined address within the Internet of a secondary server on which a copy of the sub-page is hosted.
- 10 2. A method according to claim 1 wherein the at least one signifier is a URL, and the method further comprises the steps of resolving the URL to an IP address of a secondary server on which a copy of the sub-page is hosted.
- 15 3. A method according to claim 2 wherein the resolution of the URL to an IP address takes place within a main server hosting the first web page.
- 20 4. A method according to claim 2 wherein the resolution of the URL to an IP address takes place within a DNS server.
- 25 5. A method according to claim 1 wherein the at least one signifier is an IP address of a secondary server on which a copy of the sub-page is hosted.
- 30 6. A method according to any one of the preceding claims wherein the link is an alias for a plurality of alternate links, each of which has at least one signifier identifying one of a plurality of alternate secondary servers.
7. A method according to claim 6 further comprising the steps of:
 - actuating a first alternate link;
 - determining, in accordance with at least one predetermined criterion, whether actuation of the first alternate link has resulted in establishment of connection with a secondary server identified by a signifier in the first alternate link; and
 - if no connection is established according to the first predetermined criterion, actuating a second alternate link.

8. A method according to claim 7 wherein the at least one predetermined criterion includes whether connection is established within a predetermined period of time.
- 5 9. A method according to claim 6 further comprising the step of actuating each alternate link substantially simultaneously.
- 10 10. A method according to claim 9 wherein, upon reaching a predetermined step in establishment of connection for one of the simultaneously actuated alternate links, the connection process for all other alternate links is aborted.
11. A method according to any one of the preceding claims wherein the first web page is a home page.

ESTABLISHMENT OF NETWORK CONNECTIONS

ABSTRACT

- 5 Connection of a client browser to a server hosting a sub-page in a website via a link from a principal page at a higher level in the hierarchy of pages is established by actuation of a link sent to the client with the principal page. The link includes a signifier unique to a given server on which the sub-page is hosted, so that resolution of the URL in the link need not be performed dynamically at the server side.

10

Fig. 4






Fig. 1




Fig. 2




Fig. 3





Fig. 4

Fig. 5






Fig. 6




Fig. 7

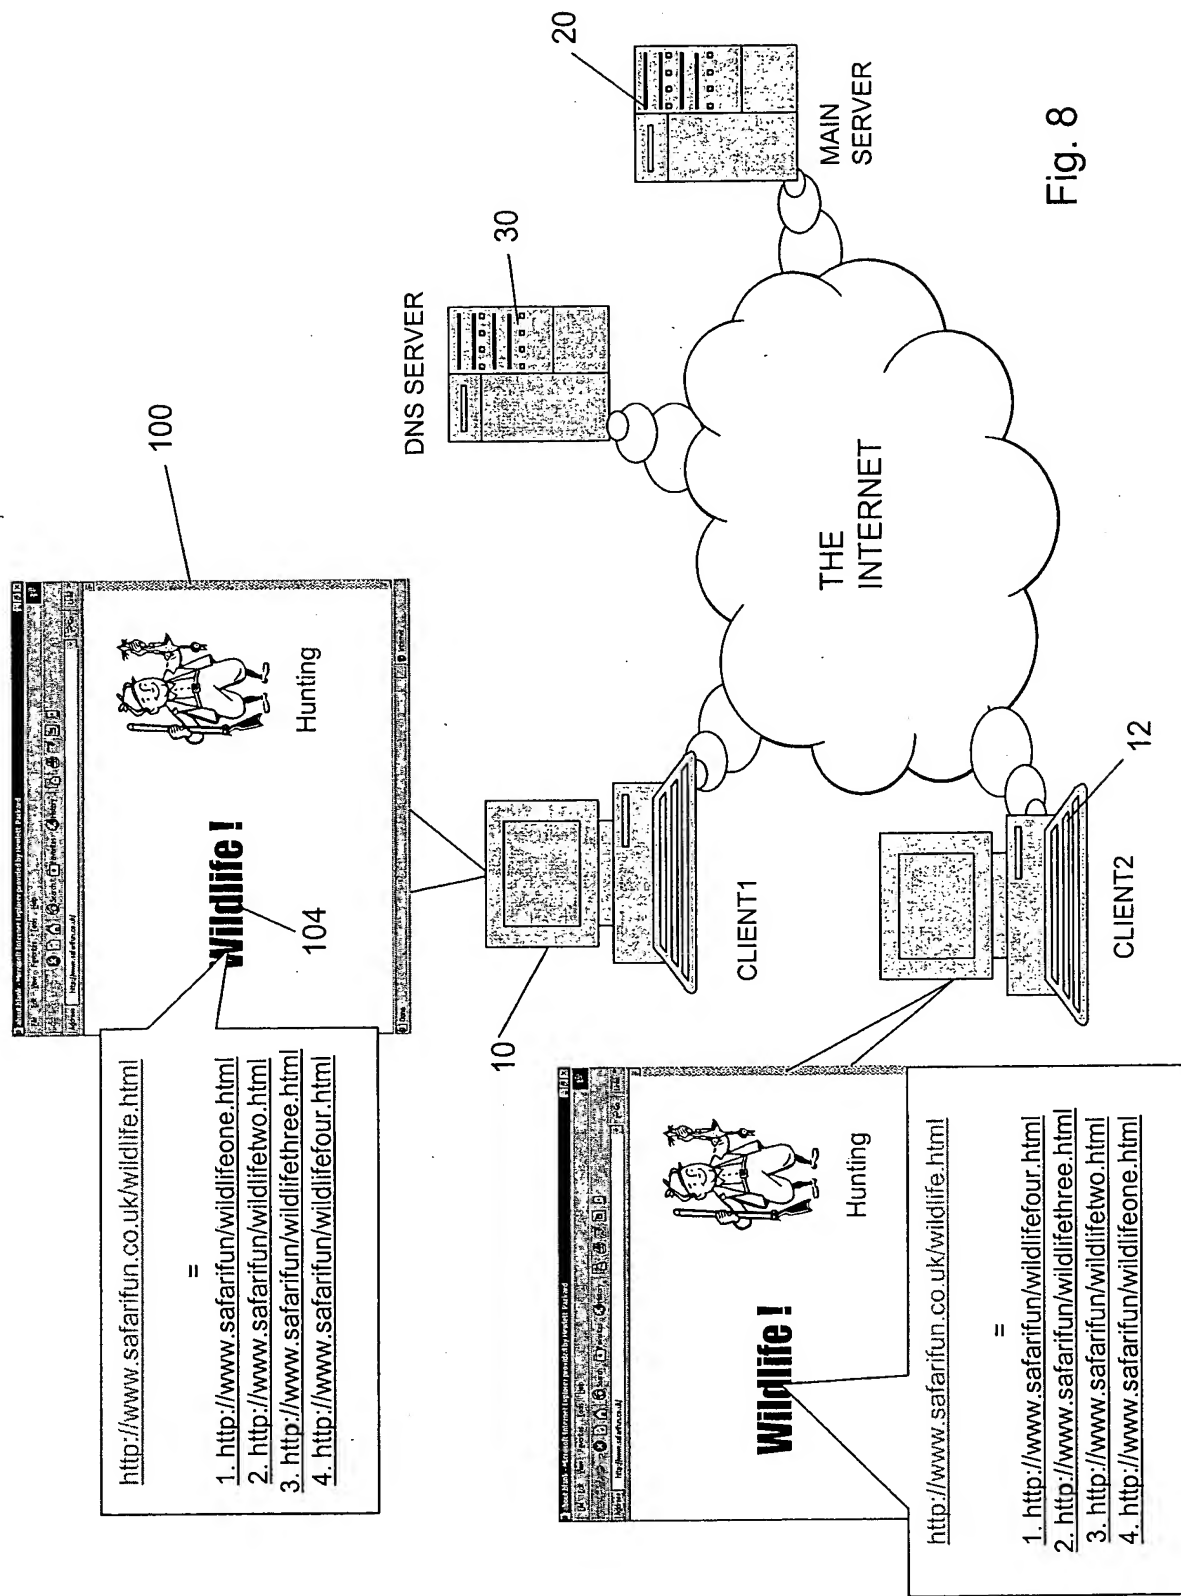


Fig. 8


```

//
// function takes a list of pairs - first member of each pair is a url
// second member is a timeout. The function loops through the list of pairs
// attempting to access the URL and waiting timeout seconds before trying
// the next. The function exits returning a 1 if any one of the urls is
// active and a -1 if none are active.
//
//

// a web page with our javascript (or a modified browser, or a plug in or java) function
// looks like this

<html>

<head>
<meta http-equiv="Content-Language" content="en-gb">
<meta http-equiv="Content-Type" content="text/html; charset=windows-1252">
<meta name="GENERATOR" content="HP-Smarlink-editor">
<meta name="ProgId" content="HPSmartlink.Editor.Document">
<title>New Page 1</title>
</head>

<body>

<p>This is a link to the <a {(href="http://www.safarifun.co.uk/wildlifeone.html",10),(href="
http://www.safarifun.co.uk/wildlifetwo.html ",10),(href=" http://www.safarifun.co.uk/wildlifethree.html
",10),(href=" http://www.safarifun.co.uk/wildlifefour.html ",10)}>wildlife</a> web page</p>

</body>

// which represents a choice of 4 sites, each of which is given 10 seconds to respond before testing the next

</html>

```

Fig. 9

